# Learning to Start for Sequence to Sequence Architecture

Qingfu Zhu*, Weinan Zhang*, Lianqiang Zhou** and Ting Liu*

*Harbin Institute of Technology,
Email:{qfzhu,wnzhang,tliu}@ir.hit.edu.cn
**Tencent Corporation,
Email:tomcatzhou@tencent.com

arXiv:1608.05554v1 [cs.CL] 19 Aug 2016

## Abstract

The sequence to sequence architecture is widely used in the response generation and neural machine translation to model the potential relationship between two sentences. It typically consists of two parts: an encoder that reads from the source sentence and a decoder that generates the target sentence word by word according to the encoder's output and the last generated word. However, it faces to the "cold start" problem when generating the first word as there is no previous word to refer. Existing work mainly use a special start symbol "</s>" to generate the first word. An obvious drawback of these work is that there is not a learnable relationship between words and the start symbol. Furthermore, it may lead to the error accumulation for decoding when the first word is incorrectly generated. In this paper, we proposed a novel approach to learning to generate the first word in the sequence to sequence architecture rather than using the start symbol. Experimental results on the task of response generation of short text conversation show that the proposed approach outperforms the state-of-the-art approach in both of the automatic and manual evaluations.

## 1 Introduction

Recently, the sequence to sequence(Seq2Seq) architecture has gained great development as a general neural network method to model the potential relationship between two sequences. For the basic Seq2Seq model, each sequence is usually modeled by RNN, and the two RNNs for the source sequence and target sequence are called encoder and decoder respectively. The encoder reads from the source sentence and do some summarize. The decoder is actually a language model that produce words according to previously predicted words conditioned with the encoder's output(usually called the context vector). This indicates that when the decoder try to predict a word, the context vector and the word predicted at previous time are two necessary inputs that requires.

So here comes the initialization question: when producing the first word by the decoder, there is no previous predicted word to be referenced to. Typically, previous work use a start symbol "</s>" to generation the first word (Sutskever et al., 2014). While it is not suitable to introduce a start symbol as the first word varies from different sentences. Concretely, there is not a learnable conditional probability of words given start symbol. Meanwhile, the process of producing the first word and generating the rest words of a sentence are different so that they should be handled respectively. To address this issue, we proposed a novel approach to learning to generate the first word. In detail, we find two factors that impact the encoding and decoding process: one is the source sequence which can be expressed using the encoder's states. The other is the representation of candidate words, of which information are all contained in the embedding matrix. We thus introduce these variables to map the representation of the source sentence into a probability distribution over the word table, pick up the maximal dimension as the final result.

The contribution of this paper is as follows:

- To the best of our knowledge, we are the first to proposed a novel approach to learning to generate the first word in Seq2Seq architecture.

- The proposed approach outperforms the state-of-the-art on the response generation of the short text conversation.

- Besides the short text conversation task, the proposed approach is a general framework which can also adapt to other Seq2Seq learning applications.

## 2 Background

From the perspective of probability, the Seq2Seq model maximize the probability of the target sequence conditioned with the source sequence during the training process, and search for a sequence that have a maximal conditioned probability given the source sequence during the predicting process. Due to that highly abstract attribute, lots of tasks such as Response Generation,Machine Translation and Question Answering can all be modeled using that architecture.

### 2.1 RNN encoder-decoder

Typically, a sequence to sequence model consists of two parts: encoder and decoder, both of which are often implemented using a family of RNN, such as GRU (Cho et al., 2014a; Chung et al., 2014) and LSTM (Hochreiter and Schmidhuber, 1997; Gers et al., 2000; Graves, 2012a), so a seq2seq model is also called RNN encoder-decoder architecture.

The encoder is a normal RNN, which reads from a sequence of words and outputs their hidden states.These states are also called annotations denoted by $H$, and for each hidden state $h_i$ at time $i$, it is computed by its previous hidden state $h_{i-1}$ and the word at time $t$:

$$h_t = f(h_{t-1}, x_t); \quad H = \{h_1, h2, h3, ..., h_T\} \tag{1}$$

Here, $T$ is the length of the source sequence, $f$ is a non-linear function. After that, the encoder computes a distributed representation using these hidden states as a summary(context vector) of the input sequence. The most simplest way is directly fetching the last one:

$$c = q(\{h_1, h2, h3, ..., h_T\}) \tag{2}$$

For the decoder, the hidden state's calculation is quite similar, the only difference is that the sequence input $x_t$ is replaced by the word predicted at last time:

$$s_t = f(s_{t-1}, y_{t-1}) \tag{3}$$

It should be noted that, the context vector $c$ is used to initialize the hidden state of decoder to make sure that the decoder was conditioned with the encoder. Based on that, (Cho et al., 2014b) add the vector $c$ as an extra input into the computation of the hidden state in decoder to make sure that every time step of the decoder can get full information of the context. In that way, the formula 3 should be updated to:

$$s_t = f(s_{t-1}, y_{t-1}, c) \tag{4}$$

Then, the word at time t can be predicted by mapping the $s_t$ to a probability over the word table using the maxout activation(Goodfellow et al., 2013).

### 2.2 Attention Mechanism in Seq2Seq

In the basic architecture of the sequence to sequence, source sequence sent to the encoder is encoded into a dense, fixed-length vector. Considering that vector may not be able to contain all the useful information of the source sequence, thus becoming a bottleneck of the model,Bahdanau et al. (2014) add the attention mechanism to improve the Seq2Seq's performance.Compared with the basic architecture, which use the last hidden state as the context vector $c$, attention mechanism gives a weight to all the annotations, then use them to calculate a weighted sum as a new context vector. It should be noted that, in that way, the

vector $c$ is distinct for every time step in the decoder, because a time-related variable was involved during the computing, so here we denote the result as vector $c_j$.

$$c_j = \sum_{i=1}^{T} \alpha_{ij} h_i \tag{5}$$

Here, $c_j$ is the context vector when we decode the $j$-th word in the decoder, and the weight $\alpha_{ij}$ for the $i$-th annotation of encoder is computed by:

$$\alpha_{ij} = a(s_{j-1}, h_i) \tag{6}$$

where $a$ is a forward neural network. Intuitively, the vector $s_{j-1}$ contains the context information of the response, so Formula 6 can be understanded as to calculate the similarity between that context and these encoder annotations, which can be also regarded as a weight.

### 2.3 Initialization in Seq2Seq Learning

Initialization is such a small detail that can be ignored easily,sometimes. However, it is an important part of the model. In the encoder RNN, a state will be used to compute the state at next time(see Eq.1), and by this way, the initial state will have an indirectly influence on all the states next. The decoder RNN share the same situation. In addition, the decoder has an extra variable that should be initialized: predicted word at last time step, because we don't have that input for the first process of generation. Typically, we set the initial hidden state of encoder to an all zero vector, and people usually use the last hidden state of the encoder to initialize the decoder's first hidden states:

$$s_0 = \sigma(W_s h_T) \tag{7}$$

where $\sigma$ is a non-linear function, $W_s$ is a trainable parameter. That is intuitively plausible because it describes the relation between the two sequences that the decoder is conditioned with the encoder. As to the previous generated word for first generation in the decoder, we manually set a start symbol to act as that role.

## 3 Learning to start

In this section, we propose a new model to accomplish Seq2Seq's the initial prediction. We think that the method using a start symbol to predict the first word is not very suitable. First, the decoder RNN is essentially a language model (Mikolov et al., 2010), which use the previous predicted words to predict a new word, from the perspective of probability, it learns a conditional probability of word that given last predicted words. While the start symbol and the first predicted word do not have such association, because most words can be put at the first position of a sequence, there is not a learnable conditional probability, so the result of taking a start symbol may cause the model prefer to predict some high frequency words, which is also observed during other conversation models using this architecture (Sordoni et al., 2015; Serban et al., 2015; Vinyals and Le, 2015). We think the reason may lie in training samples started with these words takes a higher proportion, making the decoder learn a conditional probability that given the start symbol, these words' probabilities should be higher than others. Second, we suppose the process of predicting the first word and predicting a word according to its previous word should not be treated identically, using a same method to do the two works may not be a good choice. Third, the start symbol is involved in the calculation of decoder's hidden state(see Eq. 4), so introducing a start symbol that irrelevant of a sequence and no difference between all the source sequences may indirectly influence the prediction of the rest time steps.

So we propose a new method to relive the decoder from both predicting the first word and predicting word according to the last predicted word. In our model, the first word is predicted independently from the decoder. Inspired by the initialization of the hidden state of the decoder, we use the hidden state of encoder to calculate the probability of the first word using the formula below:

$$y_0 = \sigma((\sigma(W_i c) + b_i)E + b_e) \tag{8}$$

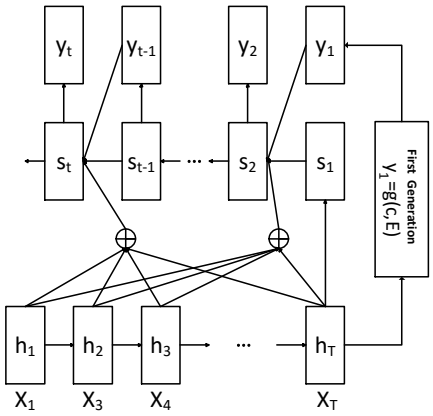


Figure 1: The LTS model architecture.

In this formula, vector $c$ is the context vector, here we directly use the last hidden state of the encoder, $W_i$ is a matrix that can be trained in the model, $E$ is the embedding matrix of the decoder, and $b_i, b_e$ are bias vectors, $\sigma$ is a no-linear activation function, so in form, the formula is equal to below if we ignore all the bias:

$$y_0 = g(c, E) \tag{9}$$

Intuitively, this formula build a tie between the context vector and the embedding matrix of the decoder. The former contains information of the source sequence, and the latter contains information of all the candidate words to be predicted in the decoder. So the $W_i$ matrix can be regarded as a similarity matrix used to compute the probability that how similar a word is to the source sequence, which indicates whether it is suitable to be predicted as the first word. Besides, by doing like this, the generation of the first word is decided only by the encoder's state. And without a start symbol's influence, the encoder's state can also be transferred to the decoder without any loss. And the rest process of prediction remains the same to the basic structure.

## 4 Experiment Settings

To verify the effectiveness, our proposed model were tested in the task of response generation of short text conversation. As a kind of neural machine architecture, a big-data is always required to get a good performance. To achieve that, a dialogue set was crawled as the training set. And to be compared with, a basic kind of Seq2Seq architecture for response generation called hybrid model proposed by Shang et al. (2015) was implemented.

### 4.1 Data

For the training process, Some one-round dialogue pairs was crawled from the Internet. For convenience, first sentence and the second sentence of one dialogue pair are denoted as post and response(Shang et al., 2015)respectively. The data set contains one million pairs, and about 35 thousands words. It should be noted that compared with the data used by Shang et al. (2015), this crawled data is a one-to-one data set, one post is corresponding to exact one response. While in the Shang et al. (2015) paper, they crawled some one-to-many data from microblog, then distributed all the responses to the its post. This is a creative way to build a big data set, while during our experiments, we found that the one-to-one data has a more rapid rate of convergence, so we created our own data set and trained models on it.Table 1 is an example of our data.

As for the test set, considering that one of our evaluation method–Bleu, which will be introduced in detail in next section, should has more than one reference for every candidate, our one-to-one data is not very suitable, so we select 100 posts and their corresponding responses in Shang et al. (2015)'s data-set to build our test set. Table 2 shows some statics of our the whole data set.

| post | response |
|---|---|
| 今天天气好差呀<br>The weather is so bad today | 雨太大了<br>The rain is too heavy |
| 每天六点多出去打篮球锻炼身体<br>Go out for exercise playing basketball at six every day | 我在打网球<br>I am playing tennis |
| 白色搭配什么颜色好<br>What color matches white well | 白色百搭呀<br>White all-match |
| 明晚我又要通宵<br>I will stay night again tomorrow night | 我陪你啦<br>I will be with you |
| 等我有空了去超市买<br>I'll go to the supermarket to buy when I was free | 超市太远了<br>The supermarket is too far away |

Table 1: Training data examples.

| Data | Data type | posts | responses |
|---|---|---|---|
| Training Data | one-to-one | 1000,000 | 1000,000 |
| Test Data | one-to-many | 1000 | 42422 |

Table 2: Data statistics

### 4.2 Models

We trained two models. The first one is a basic Seq2Seq model for dialogue called Hybird Model(denoted as HYP)(Shang et al., 2015) , the other is what we have proposed(denoted as LTS). The encoder and decoder were both implemented using the GRU (Cho et al., 2014a; Chung et al., 2014). and we set our model's parameters reference to Shang et al. (2015). The hidden size in the encoder was set to 1024. And the embedding size was set to 500, all the embedding vectors were pre-trained using the training data (Mikolov et al., 2013a; Mikolov et al., 2013b).

Besides, during the processing of training, we sent the data to the model using the mini-batch with a batch size of 100, and the RMSprop algorithm was used to update model's parameters. And we trained both models about 5 days. After that, we used the beam search algorithm to search for the N-best result of response for one paritular sentence (Graves, 2012b; Boulanger-Lewandowski et al., 2013).

## 5 Result

Until now, there is still not a uniform evaluation for response generation (Galley et al., 2015; Pietquin and Hastie, 2013; Schatzmann et al., 2005).First, we tested our model's performance using the some statistics of the the first predicted word. Second, we evaluate the complete response to see if our model can bring the Seq2Seq architecture improvement. to achieve that goal, we use two metrics: for one hand, we employed the wildly used automatic evaluation method–blue (Papineni et al., 2002) in the area of machine translation, for the other hand, we employed the human annotation method.

### 5.1 First Generated Word Evaluation

To evaluate the generation of the first word, two aspects are taken into consideration: the accuracy rate and the diversity. The dialogue pairs in the test set are denoted as test sample and reference respectively, and the sequence generated by the model is still denoted as response. As mentioned before, the test set is a one-to-many data set so each test sample corresponding to serval references. We define a set called R-set for every sample, each sample's R-set is composed by all the first words of that sample's references, during the test process, if the first word of sample's response fall into its R-set, then it will be marked as hit. And the accuracy is the ratio of hit samples over the whole test set. Furthermore, we considered such situation: some high-frequency words(derived from the training data) are so common that nearly all the R-sets consists at least one, so a sample will easily hit its R-set as well as its first generated word is such words, for example:'I'. So we further defined the accuracy without high-frequency words, denoted by $accw$-$i$,which takes such situation into consideration: if a hit word is one of top $i$ high-frequency words, then this hit will be ignored. Particularly, $accw$-0 equals the basic accuracy that do not filter any high-frequency words.

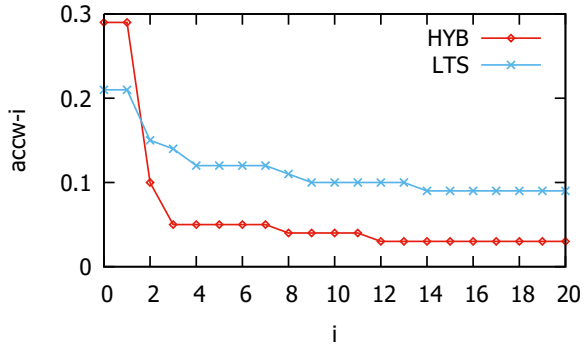


Figure 2: The accw-i metric

From the Figure 2, we can see that the LTS outperformed the HYP from the $accw$-2.We analyzed the results and find the most frequent word is a auxiliary word: "了", which seldom appear in the beginning of a sentence, so there is no change from the $accw$-0 to $accw$-1.When we ignored the second frequent word "我(I)", the performance of the HYP descends rapidly,which can be observed from the $accw$-1 to $accw$-2. while the LTS has a more stable accuracy that do not depend the easily hit high-frequency words.

Also, we evaluate the initial prediction from the perspective of diversity. In fact the Table 2 has already reflected the diversity to some degree, which our model's stable $accw$-$i$ shows that the generation of the first word is distributed fairly balance. While we still give another metric to evaluate it, we define the div-i which means the ratio of test samples whose response's first word fall into the top $i$ frequent words.According to the definition we describe, we can see that the diversity declines with increasing div-i score.

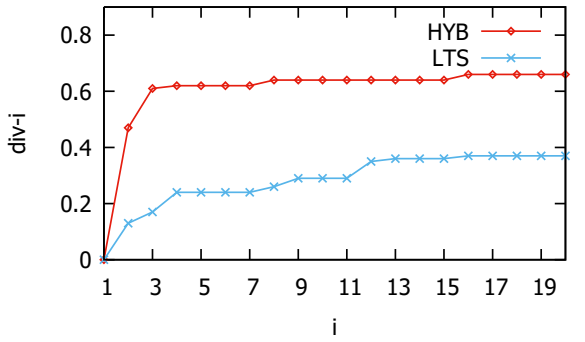


Figure 3: The div-i metric.

The Figure 3 show us that rather than concentrate on some high-frequency words, our model prefer to predict more diversity ones. Noted that in HYP there is a sharp increasement from the $div$-1 to $div$-2, which indicates that lots of the samples's first generated word is the second frequent word, which also agrees the results of $acc$-$i$.

### 5.2 Bleu Metric

We use this metric to evaluate the a model's complete response rather than the first word. which is proved to agree well with human judgement on response generation task (Sordoni et al., 2015; Li et al., 2015). And the result is given in the Table 3.

From the Table 3, we can see that the LTS performs well than the HYP in Bleu-1 to Bleu-3. Through that table, we can also see that improvement on the Bleu-1 is not as significant as other two. We analyzed this situation and got an opinion, it may because the Bleu metric calculate overlap of n-grams between response and references, compared with other n-grams, the unigram is more easily to be matched making

| | BLEU-1 | BLEU-2 | BLEU-3 |
|---|---|---|---|
| HYP | 0.5283 | 0.0553 | 0.0013 |
| LTS | **0.5303** | **0.0816** | **0.0063** |

Table 3: Bleu score.

the Bleu-1 not distinguished enough.Table 4 shows some results of the bleu evaluation, two models got similar Bleu-1 scores, while the Bleu-2 and Bleu-3 are much more strict metrics that can reflect the improvement more significantly, which also proves our analyzed mentioned before.

| No. | model | post | responds | Bleu-1 | Bleu-2 | Bleu-3 |
|---|---|---|---|---|---|---|
| 1 | HYP | 拍照中，这个造型像个娃娃一样！<br>Taking pictures, this pose looks like a little doll! | 不像男的<br>Doesn't like a man | 0.750 | 0.000 | 0.000 |
| | LTS | | 都是小娃娃<br>All are little dolls | 1.000 | 0.500 | 0.167 |
| 2 | HYP | 低调的境界<br>The state of low-key | 我也很低呀<br>I'm very low,too | 0.250 | 0.000 | 0.000 |
| | LTS | | 我很低调的<br>I am very low-key | 0.800 | 0.400 | 0.133 |

Table 4: Bleu evaluation examples.

### 5.3 Manual Evaluation

At the same time, we also tested our model adopting human annotation method. The evaluation metric is made reference to Shang (2015). We generated responses with HYP model and LTS respectively, then these responses together with their original questions are mixed up into a new file to make sure that labers can judge the result fairly. Three labers were involved to assigned a score to these responses in range of 0 to 2, and the score metric is as follows:

0:This indicates a bad response. if a response has grammar,fluency mistakes, not logic consist or relevant with the original post in semantic, it should be assigned a 0 score.

1:This means the response may not be a perfect one, but may be treated as a suitable response in some particular scenario, or it is a too general response, like "I don't know".

2:This indicates a quite appropriate response, a response can be sorted to this category only when it is free of grammar and fluency errors and is independent of scenario.

| No. | post | response | score |
|---|---|---|---|
| 1 | 咳嗽不止，有啥良策吗<br>Couldn't stop cough, any good idea? | 吃点退烧药<br>eat some antipyretics | 0 |
| 2 | 永远不要对一个外行聊你的专业<br>Never talk about your major with strangers | 我很专业的<br>I'm very professional | 0 |
| 3 | 真实的团队精神是这样子的<br>This is what real team spirit likes | 是挺厉害的<br>It's pretty powerful. | 1 |
| 4 | 哪本新华字典是你用过的<br>which XinHua dictionary is the one you used | 两本都是<br>Both of them are | 1 |
| 5 | 大熊加油，我们的家)<br>Come on for our home, Bill | 我会努力的<br>I'll try my best | 2 |

Table 5: Annotation metric examples.

Table 5 show the annotation metric in more detail. The example one conflicts with the logic consistency principle, the post said he got a cough, while the response advised the antipyretics, which is not logical relevant. The example two's response is not semantic relevant to its post, so they got a score of zero. In the example three, the response can be seen as a suitable one but it is too general, so it only got a score of one. While the example four got the same score in a different way that its response strongly depend on a particular scenario that the author of the post must has exactly two dictionaries. The last example show a suitable response that free of the questions that mentioned before.

The human annotation result was used to compute these metrics for the two model respectively: mean score, ratios of different categories. From the table 4, we can see that the LTS outperforms the HYP in

| Models | Mean Score | 0 | 1 | 2 | Agreement |
|---|---|---|---|---|---|
| HYB | 0.510 | 66.0% | 17.0% | 17.0% | 0.230 |
| LTS | **0.590** | **59.7%** | **21.7%** | **18.6%** | 0.206 |

Table 6: Annotation result.

all metrics. Besides, we evaluated the consistent of different labelers using the Fleiss' kappa, which is listed the Agreement column. We can see that both of the two models' agreement all fall into the range of 0.2-0.4, which indicates that the result is a fair agreement one.

| No. | post | HYP | LTS |
|---|---|---|---|
| 1 | 你们听说过搜狗输入法石化版吗?<br>Have you heard of of sogou pinyin's petrochemical version | 这是国服<br>This is National Server | 听说挺好的<br>I heard it's nice |
| 2 | 别管我为什么叫卢沟桥<br>Don't worry about me why called Marco Polo bridge | 那叫什么啊<br>What's it called? | 谁管你了<br>Who cares for |
| 3 | 这才是大学最大的遗憾<br>This is the greatest regret in university | 我也怕大学<br>I also afraid of university | 有啥遗憾的<br>There is nothing to regret |
| 4 | "师兄帮帮忙网站"城市开放计划。<br>"Ask Brother for Help websitde" Urban open plan | 我很开放的<br>I am very open | 很好的项目<br>It's a good prodject |
| 5 | 现在开会需要用上喇叭扩音了……<br>Now the meeting need to use the horn to be heard…… | 不需要补的<br>Don't need to fill | 没钱买喇叭<br>Don't have the money to buy a horn |

Table 7: Some results of two models.

## 6 Related Work

### 6.1 Sequence to sequence for Machine Translation

Using the sequence to sequence model, neural machine translation has already got a comparable performance to the traditional methods (Bahdanau et al., 2014). As far as we know, it was first introduced into this area by Kalchbrenner and Blunsom (2013),Sutskever et al. (2014),Cho et al. (2014b),Gao et al. (2014). Besides, Cho et al. (2014b) added the vector c as an extra input to every time step of the decoder, by doing like this, all the steps not only the first one, can get full information of the context vector. Furthermore, Bahdanau et al. (2014) proposed a novel method to calculated a weighted sum of all the annotations of the encoder. This mechanism can be regarded as a kind of attention, which means when we decode a word we chose which part of the annotations should be paid more attention to.

### 6.2 Sequence to sequence in Response Generation

General speaking, dialogue systems can be sorted into two classes (Serban et al., 2016): goal-driven represented by systems Gašić et al. (2013) and non-goal-driven systems. The neural networks methods are mainly used in the later, because a large scale of data is more easily to get in that area. Ritter et al. (2011) first combine micro-blogging data with the generative probabilistic models, then Shang et al. (2015) used this type of data on the Seq2Seq to build a short conservation machine. followed by Serban et al. (2016), who came up with the Hierarchical Nerual Network model, aiming to model the utterances and interactive structure to build a multi-round dialogue system.

At the same time, Banchs (2012) proposed methods using a different type of data,the movie dialogue. Based on that, Ameixa et al. (2014) find using the retrieal system and movie subtitles can also improvement the performance.

## 7 Conclusions

In this paper, we proposed a new approach for the sequence to sequence model to generate the first word. Proved our proposed model can bring a promotion in both the accuracy and the diversity for the first word's generation, thus improving the whole performance of the generation. Experiments in the response generation tasked verified our model's effectiveness, while rather than a method for a specific task, our proposed method is a general framework, which can also used for other tasks.

## References


[Ameixa et al.2014] David Ameixa, Luisa Coheur, Pedro Fialho, and Paulo Quaresma. 2014. Luke, i am your father: dealing with out-of-domain requests by using movies subtitles. In *International Conference on Intelligent Virtual Agents*, pages 13–21. Springer.

[Bahdanau et al.2014] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

[Banchs2012] Rafael E Banchs. 2012. Movie-dic: a movie dialogue corpus for research and development. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers-Volume 2*, pages 203–207. Association for Computational Linguistics.

[Boulanger-Lewandowski et al.2013] Nicolas Boulanger-Lewandowski, Yoshua Bengio, and Pascal Vincent. 2013. Audio chord recognition with recurrent neural networks. In *ISMIR*, pages 335–340. Citeseer.

[Cho et al.2014a] Kyunghyun Cho, Bart Van Merrienboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014a. On the properties of neural machine translation: Encoder-decoder approaches. *Computer Science*.

[Cho et al.2014b] Kyunghyun Cho, Bart Van Merrienboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014b. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *Computer Science*.

[Chung et al.2014] Junyoung Chung, Caglar Gulcehre, Kyung Hyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *Eprint Arxiv*.

[Galley et al.2015] Michel Galley, Chris Brockett, Alessandro Sordoni, Yangfeng Ji, Michael Auli, Chris Quirk, Margaret Mitchell, Jianfeng Gao, and Bill Dolan. 2015. deltableu: A discriminative metric for generation tasks with intrinsically diverse targets. *arXiv preprint arXiv:1506.06863*.

[Gao et al.2014] Jianfeng Gao, Xiaodong He, Wen-tau Yih, and Li Deng. 2014. Learning continuous phrase representations for translation modeling. In *In ACL*. Citeseer.

[Gašić et al.2013] M Gašić, Catherine Breslin, Matthew Henderson, Dongho Kim, Martin Szummer, Blaise Thomson, Pirros Tsiakoulis, and Steve Young. 2013. On-line policy optimisation of bayesian spoken dialogue systems via human interaction. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 8367–8371. IEEE.

[Gers et al.2000] Felix A Gers, Jürgen Schmidhuber, and Fred Cummins. 2000. Learning to forget: Continual prediction with lstm. *Neural computation*, 12(10):2451–2471.

[Goodfellow et al.2013] Ian J Goodfellow, David Warde-Farley, Mehdi Mirza, Aaron C Courville, and Yoshua Bengio. 2013. Maxout networks. *ICML (3)*, 28:1319–1327.

[Graves2012a] Alex Graves. 2012a. Neural networks. In *Supervised Sequence Labelling with Recurrent Neural Networks*, pages 15–35. Springer.

[Graves2012b] Alex Graves. 2012b. Sequence transduction with recurrent neural networks. *arXiv preprint arXiv:1211.3711*.

[Hochreiter and Schmidhuber1997] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.

[Kalchbrenner and Blunsom2013] Nal Kalchbrenner and Phil Blunsom. 2013. Recurrent continuous translation models. In *EMNLP*, volume 3, page 413.

[Li et al.2015] Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2015. A diversity-promoting objective function for neural conversation models. *arXiv preprint arXiv:1510.03055*.

[Mikolov et al.2010] Tomas Mikolov, Martin Karafiát, Lukas Burget, Jan Cernockỳ, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *Interspeech*, volume 2, page 3.

[Mikolov et al.2013a] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

[Mikolov et al.2013b] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.

[Papineni et al.2002] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.

[Pietquin and Hastie2013] Olivier Pietquin and Helen Hastie. 2013. A survey on metrics for the evaluation of user simulations. *The knowledge engineering review*, 28(01):59–73.

[Ritter et al.2011] Alan Ritter, Colin Cherry, and William B Dolan. 2011. Data-driven response generation in social media. In *Proceedings of the conference on empirical methods in natural language processing*, pages 583–593. Association for Computational Linguistics.

[Schatzmann et al.2005] Jost Schatzmann, Kallirroi Georgila, and Steve Young. 2005. Quantitative evaluation of user simulation techniques for spoken dialogue systems. In *6th SIGdial Workshop on DISCOURSE and DIALOGUE*.

[Serban et al.2015] Iulian V Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2015. Hierarchical neural network generative models for movie dialogues. *arXiv preprint arXiv:1507.04808*.

[Serban et al.2016] Iulian V Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI-16)*.

[Shang et al.2015] Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. Neural responding machine for short-text conversation. *Computer Science*.

[Sordoni et al.2015] Alessandro Sordoni, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015. A neural network approach to context-sensitive generation of conversational responses. *arXiv preprint arXiv:1506.06714*.

[Sutskever et al.2014] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. *Advances in Neural Information Processing Systems*, 4:3104–3112.

[Vinyals and Le2015] Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.